

CLAIMS

1. A method of executing an information retrieval query in a multiserver computing environment, comprising:

 distributing the query among each of a plurality of partial index servers in the multiserver environment;

 calculating a subset of results for each of the plurality of partial index servers; and

 merging the subset of results in one logical index server to generate a merged result.

2. A method in accordance with claim 1, further comprising receiving the query from a frontend computer system.

3. A method in accordance with claim 2, further comprising sending the merged result to the frontend computer system.

4. A method in accordance with claim 1, wherein distributing the query further includes:

 requesting N unsorted results with an offset M from one of the plurality of partial index servers, the N unsorted results having keys K1, ..., KN; and

requesting N unsorted results from each other of the plurality of partial index servers, wherein each of the N unsorted results has the same keys K_j of the respective N unsorted results requested from the one of the plurality of partial index servers, and wherein $(1 \leq j \leq N)$.

5. A method in accordance with claim 1, wherein distributing the query further includes:

requesting N sorted results with an offset M from each of the plurality of partial index servers; and

merging C results from each of the plurality of partial index servers in one logical index server, wherein C is much greater than N and represents an upper bound on the number of results that needs to be considered in order to obtain the N results required.

6. A method in accordance with claim 5, further comprising selecting a number C.

7. A method in accordance with claim 5, further comprising determining whether the value of C is sufficient.

8. A method in accordance with claim 7, wherein determining whether the value of C is sufficient further includes:

determining a sum G of the Cth aggregate values for the plurality of partial index servers;

determining a value H representing the (N + M)th aggregate value; and

if G ≥ H, resetting C to a new value.

9. A system for executing an information retrieval query in a multiserver computing environment, comprising:

a frontend computer system for receiving the query;

one or more partial index servers configured to receive the query, and to calculate a subset of results; and

a logical index server configured to merge the subset of results from each of the one or more partial index servers to generate a merged result.

10. A system in accordance with claim 9, wherein the logical index server is further configured to forward the merged result to the frontend computer system.